

CLAIM AMENDMENTS

IN THE CLAIMS

This listing of the claims will replace all prior versions, and listing, of claims in the application or previous response to office action:

1. (Previously Presented) A method for allocating application processing operations among information handling system cluster resources in response to a fail-over event, comprising:

calculating a performance ratio between a performance-related characteristic identified for a failing-over cluster node and a similar performance-related characteristic identified for a fail-over cluster node;

transforming a first calendar schedule associated with failing-over application processing operations into a second calendar schedule to be associated with failing-over application processing operations on the fail-over cluster node based at least on the calculated performance ratio;

determine whether resources on the fail-over cluster node are sufficient to support both failing-over application processing operations in accordance with the second calendar schedule and any existing fail-over cluster node application processing operations;

in response to determining that the resources of the fail-over cluster node are insufficient to support both failing-over application processing operations in accordance with the second calendar schedule and any existing fail-over cluster node application processing operations, applying a resource negotiation algorithm to the second calendar schedule associated with failing-over application processing operations;

calculating a new second calendar schedule for the fail-over node application processing operations based on results from application of the resource negotiation algorithm, the new second calendar schedule providing a reduced operating state for at least a portion of the failing-over application processing operations, but not reducing the operating state for any existing fail-over cluster node application processing operations; and

implementing the new second calendar schedule on the fail-over cluster node such that the fail-over cluster node may effect failing-over application processing operations according to the new second calendar schedule.

2-4. Cancelled.

5. (Original) The method of Claim 1, further comprising:
identifying at least one characteristic of the failing-over cluster node;
identifying at least one characteristic of the fail-over cluster node; and
calculating the performance ratio between the failing-over cluster node and the fail-over cluster node based on the identified characteristics of each node.

6. (Original) The method of Claim 1, further comprising collecting information handling system cluster node resources required by at least one application to be deployed in an information handling system cluster configuration.

7. (Original) The method of Claim 1, further comprising maintaining a knowledge-base containing information regarding one or more operational aspects of the information handling system cluster.

8. (Original) The method of Claim 7, further comprising determining whether the first calendar schedule for a selected cluster node is feasible using operational aspects of the selected cluster node available in the knowledge-base.

9. (Original) The method of Claim 1, further comprising updating an application-to-cluster node map identifying the cluster node associated with each application following the allocation of application processing operations among the information handling system resources in response to a fail-over event.

10. (Previously Presented) A system for maintaining resource availability in response to a fail-over event, comprising:

an information handling system cluster including a plurality of nodes;

at least one storage device operably coupled to the cluster; and

a program of instructions storable in a memory and executable in a processor of at least one node, the program of instructions operable to:

identify at least one performance-related characteristic of a failing node and at least one performance-related characteristic of a fail-over node;

calculate a performance ratio between the performance-related characteristic identified for the failing node and the performance-related characteristic identified for the fail-over node;

transform a processing schedule for at least one failing-over application to a new processing schedule associated with failing-over application processing on the fail-over node based at least on the calculated performance ratio;

determine whether application resources available on the fail-over node are sufficient to perform processing operations for the failing-over application in accordance with the new processing schedule and any existing fail-over application processing operations

apply a resource negotiation algorithm to at least the new processing schedule in response to a determination that the application processing resources of the fail-over node are insufficient to support both the processing schedule of the failing-over application and any existing fail-over applications;

calculate a modified processing schedule in accordance with results of the resource negotiation algorithm, the modified processing schedule defining a reduced operating state for at least a portion of the failing-over application processing operations, but not reducing the operating state for any existing fail-over cluster node application processing operations; and

implement the modified processing schedule for the failing-over application on the fail-over node.

11. (Original) The system of Claim 10, further comprising the program of instructions operable to gather node resource requirements for at least one application to be deployed in the cluster.

12. (Original) The system of Claim 11, further comprising the program of instructions operable to gather resources available on at least one node of the cluster.

13. (Original) The system of Claim 12, further comprising the program of instructions operable to verify that the resources of a selected node are sufficient to perform processing operations in accordance with the resource requirements of at least one application to be deployed on the selected node.

14. Cancelled.

15. Cancelled.

16. (Previously Presented) The system of Claim 11, further comprising the program of instructions operable to apply the resource negotiation algorithm to the new processing schedule for the failing-over application and at least one existing fail-over node processing schedule.

17. (Previously Presented) Software for allocating information handling system resources in a cluster in response to a fail-over event, the software embodied in computer readable media and when executed operable to:

- access a knowledge-base containing application resource requirements and available cluster node resources;

- calculate a performance ratio between a performance-related characteristic identified for a failing-over node and a similar performance-related characteristic identified for a fail-over node;

- develop a new processing schedule for a failing-over application on the fail-over node based at least on the calculated performance ratio;

- determine whether the new processing schedule may be supported by the fail-over node;

- apply a resource negotiation algorithm to the new processing schedule;

- generate a modified new processing schedule based on the application of the resource negotiation algorithm, the modified new processing schedule selectively reducing an operating state for the failing-over application but no applications existing on the fail-over node at the time of the fail-over event; and

- queue the failing-over application for processing on the fail-over node in accordance with the modified new processing schedule.

18. (Original) The software of Claim 17, further operable to:

- gather resource requirements for each application in the cluster selected for fail-over protection; and

- store the application resource requirements in a static data portion of the knowledge-base.

19. (Original) The software of Claim 18, further operable to:

- gather available resource information for each cluster node selected for operation as a fail-over node; and

- store the available node resource information in the static data portion of the knowledge-base.

20. (Original) The software of Claim 19 further operable to determine whether a selected node includes resources available to support a processing schedule for a selected application based on the resource requirements of the application and the available resources on the node from information maintained in the knowledge-base.

21. Cancelled.

22. (Original) The software of Claim 21, further operable to:
apply a resource negotiation algorithm to each processing schedule associated with the fail-over node;
generate new processing schedules for applications to be executed by the fail-over node; and
queue the applications to be executed by the fail-over node in accordance with resource negotiation algorithm generated processing schedules.

23. (Original) The software of Claim 17, further operable to update an application-to-node map contained in the knowledge-base.